08128/00

A1

UTILITY PATENT APPLICATION TRANSMITTAL

(New Nonprovisional Applications Under 37 CFR § 1.53(b))

Attorney Docket No. 002.0141.01

TO THE ASSISTANT COMMISSIONER FOR PATENTS:

Transmitted herewith is the patent application of ( ) application identifier or (X) first named inventor, Daniel T. Holland III, entitled System And Method For Intrusion Detection Data Collection Using A Network Protocol Stack Multiplexor, for a(n):

(X) Original Patent Application.

( ) Continuing Application (prior application not abandoned):
- ( ) Continuation ( ) Divisional ( ) Continuation-in-part (CIP) of prior application No: _____ Filed on: _____.
- ( ) A statement claiming priority under 35 USC § 120 has been added to the specification.

22895
PATENT TRADEMARK OFFICE

Enclosed are:
- (X) Specification; 24 Total Pages. (X) Drawing(s); 11 Total Sheets.
- (X) Oath or Declaration:
  - (X) A Newly Executed Combined Declaration and Power of Attorney:
    ( ) Signed. ( ) Unsigned. (X) Partially Signed.
  - ( ) A Copy from a Prior Application for Continuation/Divisional (37 CFR § 1.63(d)).
    - ( ) Incorporation by Reference. The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied, is considered as being part of the disclosure of the accompanying application and is hereby incorporated herein by reference.
    - ( ) Signed Statement Deleting Inventor(s) Named in the Prior Application. (37 CFR § 163(d)(2)).
- ( ) Power of Attorney. (X) Return Receipt Postcard.
- ( ) Associate Power of Attorney. ( ) A Check in the amount of $_____ for the Filing Fee.
- ( ) Preliminary Amendment. ( ) Information Disclosure Statement and Form PTO-1449.
- ( ) A Duplicate Copy of this Form for Processing Fee Against Deposit Account.
- ( ) A Certified Copy of Priority Documents (if foreign priority is claimed).
- ( ) Statement(s) of Status as a Small Entity.
- ( ) Statement(s) of Status as a Small Entity Filed in Prior Application, Status Still Proper and Desired.
- (X) Other: Power Of Attorney By Assignee To Exclusion Of Inventor Under 37 C.F.R. § 3.71 With Revocation Of Prior Powers, Joint Assignment, Form PTO-1595, Check in the amount of $40.00 for Recordation Fee .

CLAIMS AS FILED				
FOR	NO. FILED	NO. EXTRA	RATE	FEE
Total Claims	29	9	$18.00	$ 162.00
Independent Claims	5	2	$78.00	$ 156.00
Multiple Dependent Claims (if applicable)				$0.00
Assignment Recording Fee				$0.00
Basic Filing Fee				$690.00
			Total Filing Fee	$1,008.00

Charge $_____ to Deposit Account _____ pursuant to 37 CFR § 1.25. At any time during the pendency of this application, please charge any fees required or credit any overpayment to this Deposit Account.

Respectfully submitted,

By: _____
Patrick J.S. Inouye, Esq.
Attorney of Record, Reg. No.40297

Date: August 24, 2000

Correspondence Address:

Patrick J.S. Inouye, P.S.
816 Second Avenue P.O. Box 21808
Seattle, WA 98111-3808
Phone: (206) 381-3900
Fax: (206) 381-3999

I hereby certify that this is being deposited with the U.S. Postal Service "Express Mail Post Office to Addressee" service under 37 CFR § 1.10 on the date indicated below and is addressed to:

Assistant Commissioner for Patents
Box Patent Application
Washington, D.C. 20231

By: _____

Typed Name: Patrick J.S. Inouye, Esq

Express Mail Label No.: EL584518816US

Date of Deposit: August 24, 2000

Patent Application
Docket No. 002.0141.01
NAI Docket No. 00.029.02

SYSTEM AND METHOD FOR INTRUSION DETECTION DATA COLLECTION USING A NETWORK PROTOCOL STACK MULTIPLEXOR

Cross-Reference to Related Application

This patent application claims priority under 35 U.S.C. § 119(e) to provisional patent application Serial No. 60/182,842, filed February 16, 2000, the disclosure of which is incorporated herein by reference.

Field of the Invention

The present invention relates in general to network intrusion detection data collection and, in particular, to a system and method for intrusion detection data collection using a network protocol stack multiplexor.

Background of the Invention

Enterprise computing environments typically consist of host computer systems, individual workstations, and network resources interconnected over intranetworks internal to the organization. These intranetworks, also known as local area networks, make legacy databases and information resources widely available for access and data exchange. These systems can also be interconnected to wide area networks, including public information internetworks, such as the Internet, to enable internal users access to remote data exchange and

computational resources and to allow outside users access to select internal resources for completing limited transactions or data transfer.

Unfortunately, enterprise computing environments are also susceptible to security compromises. A minority of surreptitious users, colloquially termed, "hackers," abuse computer interconnectivity by attempting to defeat security measures and intrude into non-public computer resources without authorization. Hackers pose an on-going concern for system administrators charged with safeguarding data integrity and security.

Hackers often take advantage of flaws and limitations inherent to network architectures. For instance, most internetworks and intranetworks are based on a layered network model employing a stack of standardized protocol layers. The most widely adopted network model is the Transmission Control Protocol/Internet Protocol (TCP/IP) suite, such as described in W.R. Stevens, "TCP/IP Illustrated," Vol. 1, Ch. 1 et seq., Addison-Wesley (1994), the disclosure of which is incorporated herein by reference. Computers and network resources using the TCP/IP suite implement hierarchical protocol stacks which, at minimum, include link and network layers. End-to-end devices, such as workstations and servers, further include transport and application layers.

The layering and variability of implementation in TCP/IP suites expose numerous opportunities for network compromise and exploitation by hackers. Consequently, most networks employ some form of firewall or intrusion detection system as a first line of defense against hackers. Firewalls employ packet filtering, stateful packet inspection and application proxies while intrusion detection systems typically perform signature or statistical intrusion detection. Both of these forms of security require continuous access to network traffic.

Network packet filters present one prior art solution to providing network traffic to intrusion detection systems and some forms of firewall, such as described in W.R. Stevens, "TCP/IP Illustrated," Vol. 1, App. A, Addison-Wesley (1994), the disclosure of which is incorporated herein by reference. Packet filters capture and filter data packets obtained from a network interface that has been

placed into promiscuous mode, typically by retrieving a copy from the network interface driver. Packet filters, however, suffer from several drawbacks. First, current packet filters are inherently bandwidth limited and cannot scale beyond approximately 10-20 Mbps of traffic. Packet filters also consume computational resources, including memory and processing cycles. Finally, receiving intrusion detection systems and firewalls must demultiplex raw packet traffic retrieved by packet filters into individual data packets corresponding to the individual protocol layers. The demultiplexing consumes further computational resources, duplicates work performed by the protocol stack, and introduces the potential for errors.

Therefore, there is a need for a scaleable solution to providing packet traffic for network intrusion detection and analysis. Preferably, such a solution would avoid duplication of protocol stack functionality and computational resource waste.

Summary of the Invention

The present invention provides a system and method for dynamically collecting data for use in intrusion detection directly from the network protocol stack. A stack multiplexor introduces a set of shims at select points in the data flow of traffic through the protocol stack. The shims are introduced by redirecting driver entry points in a module switch table. Copies of message blocks referring to the collected data are forwarded to an analysis module for intrusion detection and analysis.

An embodiment of the present invention is a system and method for intrusion detection data collection using a protocol stack multiplexor. A hierarchical protocol stack is defined within kernel memory space. The protocol stack includes a plurality of communicatively interfaced protocol layers. Each such protocol layer includes one or more procedures for processing data packets. A data frame is processed through the protocol stack. The data frame includes a plurality of recursively encapsulated data packets which are each encoded with a protocol recognized by one of the protocol layers. Data is collected directly from the protocol stack from at least one of the processed data packets using a protocol

stack multiplexor. Redirected references interface directly into at least one such protocol layer to the data packet processing procedures included within the at least one such protocol layer. A logical reference to the processed data packets is obtained from the interfaced protocol layer. The logical reference refers to a memory block in the kernel memory space within which the processed data packets are stored. The logical reference is provided to an intrusion detection analyzer executing within user memory space.

A further embodiment of the present invention is a system and method for detecting network intrusions using a protocol stack multiplexor. A network protocol stack includes a plurality of hierarchically structured protocol layers. Each such protocol layer includes a read queue and a write queue for staging transitory data packets and a set of procedures for processing the transitory data packets in accordance with the associated protocol. A protocol stack multiplexor is interfaced directly to at least one such protocol layer through a set of redirected pointers to the processing procedures of the interfaced protocol layer. A data packet collector references at least one of the read queue and the write queue for the associated protocol layer. A data packet exchanger communicates a memory reference to each transitory data packet from the referenced at least one of the read queue and the write queue for the associated protocol layer. An analysis module receives the communicated memory reference and performs intrusion detection based thereon.

Still other embodiments of the present invention will become readily apparent to those skilled in the art from the following detailed description, wherein is described embodiments of the invention by way of illustrating the best mode contemplated for carrying out the invention. As will be realized, the invention is capable of other and different embodiments and its several details are capable of modifications in various obvious respects, all without departing from the spirit and the scope of the present invention. Accordingly, the drawings and detailed description are to be regarded as illustrative in nature and not as restrictive.

Brief Description of the Drawings

FIGURE 1 is a block diagram showing a distributed computing environment, including a system for intrusion detection data collection using a network protocol stack multiplexor, in accordance with the present invention.

FIGURE 2 is a block diagram of a prior art system for intrusion data collection.

FIGURE 3 is block diagram of a system for intrusion detection data collection using a network protocol stack multiplexor.

FIGURE 4 is a block diagram of a Transmission Control Protocol/Internet Protocol-compliant (TCP/IP) network protocol stack implementation.

FIGURE 5 is a flow diagram of a method for intrusion detection data collection using a network protocol stack multiplexor in accordance with the present invention.

FIGURE 6 is a flow diagram of a routine for initializing a stack multiplexor for use in the method of FIGURE 5.

FIGURE 7 is a flow diagram of a routine for collecting data for use in the method of FIGURE 5.

FIGURE 8 is a flow diagram of a routine for collecting raw data frames for use in the routine of FIGURE 7.

FIGURE 9 is a flow diagram of a routine for collecting IP datagrams for use in the routine of FIGURE 7.

FIGURE 10 is a flow diagram of a routine for collecting TCP-processed data packets for use in the routine of FIGURE 7.

FIGURE 11 is a flow diagram of a routine for collecting UDP datagrams for use in the routine of FIGURE 7.

Detailed Description

FIGURE 1 is a block diagram showing a distributed computing environment 10, including a system for intrusion detection data collection using a network protocol stack multiplexor 21, in accordance with the present invention. The environment 10 includes a intranetwork 13 interconnected with an

internetwork 14, such as the Internet. The intranetwork 13 includes a local server 12 with a plurality of clients 11 and similar network resources (not shown). The intranetwork 13 is interconnected to a remote server 16 via the internetwork 14 and both the remote server 16 and the intranetwork 13 are interfaced to the internetwork 14 via routers 15. Other network topologies and configurations are feasible.

The intranetwork 14 also includes several forms of intrusion detection, including a firewall 17, a network intrusion detection system (IDS) 18, a set of host IDSs 19, and a hybrid IDS 20. The firewall 17 prevents unauthorized access to the intranetwork using packet filtering, stateful packet inspection, and application proxies. The network IDS 18, host IDSs 19, and hybrid IDS 20 all collect and analyze a traffic stream to detect any attempts or actual compromises of network or system security. The network IDS 18 focuses on all traffic entering the intranetwork 18 and analyzes that traffic using signature-based and statistical-based intrusion detection techniques. Each host IDS 19 focuses on activities within their respective client 11 through internal security auditing mechanisms. The hybrid IDS 20 focuses on incoming traffic as well as internal activities and can include a protocol stack multiplexor 21 (MUX) for collecting data for use in intrusion detection, as further described below beginning with reference to FIGURE 3. An exemplary IDS is the CyberCop Monitor product, licensed by Network Associates, Inc., Santa Clara, California. Firewalls, IDSs, and related network security concerns are described in "Next Generation Intrusion Detection in High Speed Networks," Network Associates, Inc. (1998), the disclosure of which is incorporated herein by reference.

The individual computer systems, including clients 11, server 12, and remote server 16, are general purpose, programmed digital computing devices consisting of a central processing unit (CPU), random access memory (RAM), non-volatile secondary storage, such as a hard drive or CD ROM drive, network interfaces, and peripheral devices, including user interfacing means, such as a keyboard and display. Program code, including software programs, and data are

loaded into the RAM for execution and processing by the CPU and results are generated for display, output, transmittal, or storage.

FIGURE 2 is a block diagram of a prior art system 30 for intrusion data collection. By way of example, the system 30 is a Transmission Control Protocol/Internet Protocol-compliant (TCP/IP) computing environment, such as described in W.R. Stevens, "TCP/IP Illustrated," Vol. 1, Ch. 1 et seq., Addison-Wesley (1994), the disclosure of which is incorporated herein by reference. However, the present discussion can equally be applied to other layered network architectures, including those based on the ISO/OSI model. A client 11 (shown in FIGURE 1) is physically interconnected to an intranetwork 13 (or internetwork 14) via a network interface controller (NIC) 31. Incoming data frames are processed through an internet protocol (IP) stack 33 for eventual delivery to host applications 40. Similarly, outgoing data packets originating from the host applications 40 are processed through the IP stack 33 for eventual transmission over the intranetwork 13. A C2 auditing system provides host-based security by monitoring system-level activities. A host collector 35 receives the monitoring data which is reported to an analysis module 36 for intrusion analysis and detection.

A packet filter 37 collects all network traffic transiting through the NIC 31. The NIC 31 is left in standard mode, that is, a mode which copies out all network traffic destined for the media access control (MAC) address of that NIC 31 only and includes, but is not limited to, specified ports, inbound and outbound traffic, and specific protocols. The packet filter 37 captures and filters the data frames. A stream and packet processing module 38 demultiplexes the filtered data frames into individual frames, datagrams, and packets in accordance with the network protocols supported by the IP stack 33. In effect, the stream and packet processing module duplicates the functionality of the IP stack 33 by reassembling raw data frames into properly formatted, higher protocol data packets. These data packets are collected by a network collector 39 for use by the analysis module 36.

Both the IP stack 33 and C2 auditing system operate in kernel memory space 32 while the remaining components operate in user memory space. The kernel memory space 32 is privileged memory space used for and controlled exclusively by the operating system. Transitioning data values to and from the kernel memory space 32 involves a context switch and incurs a performance penalty.

As a hardware device, the NIC 33 is outside the kernel memory space 32 but the actual copying of the network traffic from the NIC 33 to the packet filter 37 is performed by a network driver (not shown) also operating in the kernel memory space 32. Consequently, the copying of each data frame is computationally expensive due to the context switch and sheer volume of data copied. Similarly, the demultiplexing of raw data by the stream and packet processing module 38 duplicates the work performed by the IP stack 33 and introduces the potential for erroneously reassembled packets. These shortcomings can be exploited by a would-be network intruder and introduces problems when trying to accurately detect certain types of attacks.

FIGURE 3 is block diagram of a system for intrusion detection data collection 50 using a network protocol stack multiplexor 62. Raw network traffic transits to and from the intranetwork 13 (or internetwork 14) through the NIC 51 and is processed though the IP stack 52. The C2 auditing system 57 provides host-based security by monitoring system-level activities. The host collector 60 receives the monitoring data which is reported to the analysis module 61. The IP stack 52 and C2 auditing system both operate in kernel memory space 68.

In the described embodiment, the IP stack 52 is implemented as a Streams-based stack for use in a Unix System V, Release 4, (SVR4) compliant operating environment. The device end of the IP stack 52 at the juncture between software and hardware is referred to as the *driver end.* The user end of the IP stack 52 at the juncture between user memory space and kernel memory space is referred to as the *stream head.* The IP stack 52 is structured into hierarchical protocol layers which include internet protocol (IP) layer 53, transmission control

protocol (TCP) layer 54, and user datagram protocol (UDP) layer 55, plus other routines for processing other protocols as the remaining implementation 56. Incoming packets are forwarded to and outgoing packets originate from a set of host applications 59.

In addition to the NIC 51, select individual protocol layers between the driver end and the stream head, including IP layer 53, TCP layer 54, and UDP layer 55, are "shimmed" into the protocol stack multiplexor 62 at key data flow points, as further described below with reference to FIGURE 4. Copies of the message blocks for each processed data packet, rather than copies of the data packets themselves, are received by the stack multiplexor 62 for raw data (RAW_DATA) 67, IP data (IP_DATA) 66, UDP data (UDP_DATA) 65, and TCP data (TCP_DATA) 64. No packet filtering or other processing is performed. A network capture module 63 collects the message blocks for use by the analysis module 61.

A module switch table (MST) 58 is also maintained in the kernel memory space 58. Each protocol layer is implemented as a stream driver. This table stores the entry points to the services that each stream driver provides. Each service is itself a procedure used for data packet processing. In the described embodiment, there are six main entry points, as follows:

Open	Called when a connection is initiated to the driver.
Close	Called when a connection is closed.
Readput	Called when data needs to be placed in the Read Queue.
Writeput	Called when data needs to be placed in the Write Queue.
ReadService	Called when data cannot be put into the Read Queue and for deferred processing of data packets traveling upstream from the Driver End.
WriteService	Called when data cannot be put into the Write Queue and for deferred processing of data packets traveling downstream from the Stream Head.

Other entry points and data packet processing procedures, including operating system dependent entry points, are feasible.

Each module in the stack multiplexor 62 is a computer program or module written as source code in a conventional programming language, such as the C++ programming languages, and is presented for execution by the CPU as object or byte code, as is known in the art. The various implementations of the source code and object and byte codes can be held on a computer-readable storage medium or embodied on a transmission medium in a carrier wave.

The stack multiplexor 62 operates in accordance with a sequence of process steps, as further described below beginning with reference to FIGURE 5.

FIGURE 4 is a block diagram of a Transmission Control Protocol/Internet Protocol-compliant (TCP/IP) network protocol stack implementation 80. The protocol layers are categorized into four layers, link layer 81, network layer 82, transport layer 83, and application layer 84. The link layer 81, network layer 82, and transport layer 83 operate in kernel memory space 85 while the application layer 84 operates in user data space 86.

In the described embodiment, data is collected from four protocol layer implementations using "shims" inserted at key locations in the data traffic stream. Although described with reference to upstream traffic flow from the driver end to the stream head, the present invention can equally apply to downstream traffic flow. Thus, raw incoming data frames 92 are tapped from the link layer 81 via a network interface controller 87. IP datagrams 95 are tapped from the network layer 82 via the IP layer 88. Finally, data packets and UDP datagrams are tapped from the transport layer 83 via the TCP layer 89 and UDP layer 90, respectively. TCP segments 98 and processed UDP datagrams 105 are ignored.

Using the Streams-based approach, each protocol layer implementation includes a pair of read queues 93, 96, 99, 102 and write queues 94, 97, 100, 103 for the NIC 87, IP layer 88, TCP layer 99, and UDP layer 90, respectively. The location of the shim depends upon the nature of the data being collected. Raw, IP, and UDP data are packed-based, so traffic originating from the NIC 87, IP layer

88, and UDP layer 90 can be collected directly from the respective read queues 93, 96, 102. However, TCP data is connection-based, so traffic must be collected after the IP layer 88 has completed processing of incoming TCP segments 98. A separate module (not shown) including a separate pair of read and write queues is introduced upstream from the TCP layer 99 and data packets 104 are collected from this upstream read queue.

FIGURE 5 is a flow diagram of a method 120 for intrusion detection data collection using a network protocol stack multiplexor 62 (shown in FIGURE 3) in accordance with the present invention. The method 120 operates in two phases. During the first phase (blocks 121-122), initialization, the IP stack 52 is initialized (block 121) by registering the driver entry points in the module switch table 58 and starting each driver. In addition, the protocol stack multiplexor 62 is initialized (block 122) to redirect select driver entry points, as further described below with reference to FIGURE 6.

During the second phase (blocks 123-126), operation, data packets are processed in two threads of execution (blocks 124 and 125). In a first thread, data frames traveling upstream from the Driver End are processed through the IP stack 52 (block 124). In a second thread, data in the form of memory block references is collected directly from the IP stack 52 (block 125), as further described below with reference to FIGURE 5. The second phase (blocks 123-126) continues indefinitely until the routine is terminated.

In the described embodiment, data is collected from data frames traveling upstream from the Driver End, but the present invention can equally apply to data packets traveling downstream from the Stream Head.

FIGURE 6 is a flow diagram of a routine 140 for initializing a stack multiplexor 62 for use in the method of FIGURE 5. The purpose of this routine is to redirect the entry points for select protocol layers in the IP stack 52. First, the module switch table 58 is copied (block 141) from the kernel memory space 68. Next, the driver entry points for select protocol layers (block 142), specifically,

the link layer 81, network layer 82, and transport layer 83 (shown in FIGURE 4) are determined. The driver entry points are then redirected as follows.

The driver entries in the module switch table 58 for the NIC 87, IP layer 88, and UDP layer 90 are selectively redirected to the stack multiplexor 62 (block 143). Both link layer 81 and network layer 82 protocols implement standardized Data Link Provider Interfaces (DLPIs). These interfaces allow network traffic to be directly tapped from the NIC 87 and IP layer 88. UDP is a packed-based protocol, so UDP datagrams 101 are captured by redirecting the *Readput* service routine for the UDP layer 90.

The driver entries for the TCP layer 89 are redirected to the stack multiplexor 62 (block 144) by introducing a separate data collection module upstream from the TCP layer 89. This data collection module includes a separate pair of read and write queues. The driver entries in the module switch table 58 are redirected to this data collection module and memory block references to the packets 104 processed by the TCP layer 89 are captured prior to forwarding the data packets 104 to the applications layer 91.

In the described embodiment, two kernel service routines, *attach* and *detach*, are used to redirect the driver entry points. When a driver is loaded, the *attach* service routine is called to publish the entry points in the module switch table 58 and to register the services to which the driver is to be linked. Similarly, when a driver is unloaded, the *detach* routine is called to unlink the driver from the registered services and to remove the entry points from the module switch table 58. The shims are created by saving existing entry points in the module switch table 58 and separately reattaching them within the stack multiplexor 62.

Upon completion of stack multiplexor 58 initialization, the routine returns.

FIGURE 7 is a flow diagram of a routine 150 for collecting data for use in the method of FIGURE 5. The purpose of this routine is to collect the various types of data from the individual protocol layers. Thus, depending upon the type of data (block 151), the appropriate routine is dispatched to collect raw data (block 152), IP data (block 153), TCP data (block 154), and UDP data (block

155), as further described below with respect to FIGURES 8, 9, 10, and 11, respectively. If further data remains to be collected (block 156), the routine continues dispatch. Otherwise, the routine returns.

FIGURE 8 is a flow diagram of a routine 160 for collecting raw data frames for use in the routine of FIGURE 7. The purpose of this routine is to collect raw data frames 92 from the read queue 93 of the NIC 87. If a new data frame 92 has arrived in the read queue 93 (block 161), the message block pointer for the new data frame 92 is copied and the reference counter is incremented (block 162). The message block pointer is then forwarded to the analysis module 61 (block 163). If further data frames 92 remain (block 164), the routine continues collections. Otherwise, the routine returns.

FIGURE 9 is a flow diagram of a routine 170 for collecting IP datagrams for use in the routine of FIGURE 7. The purpose of this routine is to collect IP datagrams 95 from the read queue 96 of the IP layer 88. If a new IP datagram 95 has arrived in the read queue 96 (block 171), the message block pointer for the new IP datagram 95 is copied and the reference counter is incremented (block 172). The message block pointer is then forwarded to the analysis module 61 (block 173). If further IP datagrams 95 remain (block 174), the routine continues collections. Otherwise, the routine returns.

FIGURE 10 is a flow diagram of a routine 180 for TCP-processed data packets for use in the routine of FIGURE 7. The purpose of this routine is to collect TCP-processed data packets 104 from the read queue of a data collection layer introduced upstream from the TCP layer 89. If a new data packet 104 has arrived in the upstream read queue (block 181), the message block pointer for the new data packet 104 is copied and the reference counter is incremented (block 182). The message block pointer is then forwarded to the analysis module 61 (block 183). Similarly, if a new data packet 104 has arrived in the upstream write queue (block 184), the new data packet is forwarded to the TCP layer 89 (block 185). If further data packets 104 remain (block 186), the routine continues collections. Otherwise, the routine returns.

FIGURE 11 is a flow diagram of a routine 190 for collecting UDP datagrams for use in the routine of FIGURE 7. The purpose of this routine is to collect UDP datagrams 101 from the read queue 102 of the UDP layer 90. If a new UDP datagram 101 has arrived in the read queue 102 (block 191), the message block pointer for the new UDP datagram 101 is copied and the reference counter is incremented (block 192). The message block pointer is then forwarded to the analysis module 61 (block 193). If further UDP datagrams 101 remain (block 194), the routine continues collections. Otherwise, the routine returns.

While the invention has been particularly shown and described as referenced to the embodiments thereof, those skilled in the art will understand that the foregoing and other changes in form and detail may be made therein without departing from the spirit and scope of the invention.

What is claimed is:

1. A system for intrusion detection data collection using a protocol stack multiplexor, comprising:

a hierarchical protocol stack defined within kernel memory space and comprising a plurality of communicatively interfaced protocol layers, each such protocol layer comprising one or more procedures for processing data packets;

a data frame processed through the protocol stack, the data frame comprising a plurality of recursively encapsulated data packets which are each encoded with a protocol recognized by one of the protocol layers; and

a protocol stack multiplexor collecting data directly from the protocol stack from at least one of the processed data packets, comprising:

an interface interfacing directly into at least one such protocol layer through redirected references to the data packet processing procedures comprised within the at least one such protocol layer; and

a logical reference to the processed data packets obtained from the interfaced protocol layer, the logical reference referring to a memory block in the kernel memory space within which the processed data packets are stored and provided to an intrusion detection analyzer executing within user memory space.

2. A system according to Claim 1, further comprising:

a network hardware interface in a link protocol layer logically located at a device end of the protocol stack;

an application software interface in a transport protocol layer logically located at a user end of the protocol stack; and

the protocol stack multiplexor tapping the collected data from the protocol stack between and through the link protocol layer and the transport protocol layer.

3. A system according to Claim 2, wherein the protocol stack comprises a Transmission Control Protocol/Internet Protocol-compliant (TCP/IP) protocol stack.

4. A system according to Claim 1, further comprising:

a read queue associated with each protocol layer storing incoming data frames;

a write queue associated with each protocol layer storing outgoing data frame; and

the protocol stack multiplexor retrieving the logical reference to the processed data packets from at least one of the read queue and the write queue.

5. A system according to Claim 1, further comprising:

a module switch table in the kernel memory space storing the references to the data packet processing procedures comprised within the at least one such protocol layer; and

an initialization module in the protocol stack multiplexor replacing select procedure references in the module switch table with references to data collection procedures in the protocol stack multiplexor.

6. A system according to Claim 5, wherein one such protocol layer comprises a Transmission Control Protocol-compliant (TCP) protocol layer, further comprising:

the initialization module augmenting the procedure references in the module switch table for the procedures for processing data frames for the TCP protocol layer with references to TCP data collection procedures in the protocol stack multiplexor.

7. A system according to Claim 5, wherein one such protocol layer comprises a User Datagram Protocol-compliant (UDP) protocol layer, further comprising:

the initialization module replacing the procedure references in the module switch table for the procedures for processing incoming data frames for the UDP protocol layer with references to UDP data collection procedures in the protocol stack multiplexor.

8. A method for intrusion detection data collection using a protocol stack multiplexor, comprising:

defining a hierarchical protocol stack within kernel memory space and comprising a plurality of communicatively interfaced protocol layers, each such protocol layer comprising one or more procedures for processing data packets;

processing a data frame through the protocol stack, the data frame comprising a plurality of recursively encapsulated data packets which are each encoded with a protocol recognized by one of the protocol layers; and

collecting data directly from the protocol stack from at least one of the processed data packets using a protocol stack multiplexor, comprising:

interfacing directly into at least one such protocol layer through redirected references to the data packet processing procedures comprised within the at least one such protocol layer;

obtaining a logical reference to the processed data packets from the interfaced protocol layer, the logical reference referring to a memory block in the kernel memory space within which the processed data packets are stored; and

providing the logical reference to an intrusion detection analyzer executing within user memory space.

9. A method according to Claim 8, further comprising:

providing a network hardware interface in a link protocol layer logically located at a device end of the protocol stack;

providing an application software interface in a transport protocol layer logically located at a user end of the protocol stack; and

tapping the collected data from the protocol stack between and through the link protocol layer and the transport protocol layer.

10. A method according to Claim 9, wherein the protocol stack comprises a Transmission Control Protocol/Internet Protocol-compliant (TCP/IP) protocol stack.

11. A method according to Claim 8, further comprising:

storing incoming data frames in a read queue associated with each protocol layer;

storing outgoing data frame in a write queue associated with each protocol layer; and

retrieving the logical reference to the processed data packets from at least one of the read queue and the write queue.

12. A method according to Claim 8, further comprising:

storing the references to the data packet processing procedures comprised within the at least one such protocol layer in a module switch table in the kernel memory space; and

replacing select procedure references in the module switch table with references to data collection procedures in the protocol stack multiplexor.

13. A method according to Claim 12, wherein one such protocol layer comprises a Transmission Control Protocol-compliant (TCP) protocol layer, further comprising:

augmenting the procedure references in the module switch table for the procedures for processing data frames for the TCP protocol layer with references to TCP data collection procedures in the protocol stack multiplexor.

14. A method according to Claim 12, wherein one such protocol layer comprises a User Datagram Protocol-compliant (UDP) protocol layer, further comprising:

replacing the procedure references in the module switch table for the procedures for processing incoming data frames for the UDP protocol layer with references to UDP data collection procedures in the protocol stack multiplexor.

15. A storage medium for intrusion detection data collection using a protocol stack multiplexor, comprising:

defining a hierarchical protocol stack within kernel memory space and comprising a plurality of communicatively interfaced protocol layers, each such protocol layer comprising one or more procedures for processing data packets;

processing a data frame through the protocol stack, the data frame comprising a plurality of recursively encapsulated data packets which are each encoded with a protocol recognized by one of the protocol layers; and

collecting data directly from the protocol stack from at least one of the processed data packets using a protocol stack multiplexor, comprising:

interfacing directly into at least one such protocol layer through redirected references to the data packet processing procedures comprised within the at least one such protocol layer;

obtaining a logical reference to the processed data packets from the interfaced protocol layer, the logical reference referring to a memory block in the kernel memory space within which the processed data packets are stored; and

providing the logical reference to an intrusion detection analyzer executing within user memory space.

16. A storage medium according to Claim 15, further comprising:

providing a network hardware interface in a link protocol layer logically located at a device end of the protocol stack;

providing an application software interface in a transport protocol layer logically located at a user end of the protocol stack; and

tapping the collected data from the protocol stack between and through the link protocol layer and the transport protocol layer.

17. A storage medium according to Claim 15, further comprising:

storing incoming data frames in a read queue associated with each protocol layer;

storing outgoing data frame in a write queue associated with each protocol layer; and

retrieving the logical reference to the processed data packets from at least one of the read queue and the write queue.

18. A storage medium according to Claim 15, further comprising:

storing the references to the data packet processing procedures comprised within the at least one such protocol layer in a module switch table in the kernel memory space; and

replacing select procedure references in the module switch table with references to data collection procedures in the protocol stack multiplexor.

19. A storage medium according to Claim 18, wherein one such protocol layer comprises a Transmission Control Protocol-compliant (TCP) protocol layer and a further such protocol layer comprises a User Datagram Protocol-compliant (UDP) protocol layer, further comprising:

augmenting the procedure references in the module switch table for the procedures for processing data frames for the TCP protocol layer with references to TCP data collection procedures in the protocol stack multiplexor; and

replacing the procedure references in the module switch table for the procedures for processing incoming data frames for the UDP protocol layer with references to UDP data collection procedures in the protocol stack multiplexor.

20. A system for detecting network intrusions using a protocol stack multiplexor, comprising:

a network protocol stack comprising a plurality of hierarchically structured protocol layers, each such protocol layer comprising a read queue and a write queue for staging transitory data packets and a set of procedures for processing the transitory data packets in accordance with the associated protocol;

a protocol stack multiplexor interfaced directly to at least one such protocol layer through a set of redirected pointers to the processing procedures of the interfaced protocol layer, further comprising:

a data packet collector referencing at least one of the read queue and the write queue for the associated protocol layer; and

a data packet exchanger communicating a memory reference to each transitory data packet from the referenced at least one of the read queue and the write queue for the associated protocol layer; and

an analysis module receiving the communicated memory reference and performing intrusion detection based thereon.

21. A system according to Claim 20, further comprising:

a module switch table storing a set of pointers to the processing procedures of the interfaced protocol layer; and

an initialization module selectively redirecting the set of pointers to a set of data collection procedures comprised in the protocol stack multiplexor.

22. A system according to Claim 21, further comprising:

a one-way shim redirecting the set of pointers for processing the transitory data packets for one of the read queue and the write queue for the associated protocol layer.

23. A system according to Claim 21, further comprising:

a two-way shim redirecting the set of pointers for processing the transitory data packets for both the read queue and the write queue for the associated protocol layer.

24. A system according to Claim 20, wherein the network protocol stack is a TCP/IP-compliant protocol stack, further comprising:

a set of TCP/IP-compliant protocol layers, selected from the group comprising at least:

a data link protocol layer;

an Internet (IP) protocol layer;

an Transmission Control Protocol (TCP) layer; and

a User Datagram Protocol (UDP) layer.

25. A method for detecting network intrusions using a protocol stack multiplexor, comprising:

executing a network protocol stack comprising a plurality of hierarchically structured protocol layers, each such protocol layer comprising a read queue and a write queue for staging transitory data packets and a set of procedures for processing the transitory data packets in accordance with the associated protocol;

interfacing a protocol stack multiplexor directly to at least one such protocol layer through a set of redirected pointers to the processing procedures of the interfaced protocol layer, further comprising:

referencing at least one of the read queue and the write queue for the associated protocol layer; and

communicating a memory reference to each transitory data packet from the referenced at least one of the read queue and the write queue for the associated protocol layer; and

receiving the communicated memory reference into an analysis module and performing intrusion detection based thereon.

26. A method according to Claim 25, further comprising:

storing a set of pointers to the processing procedures of the interfaced protocol layer into a module switch table; and

selectively redirecting the set of pointers to a set of data collection procedures comprised in the protocol stack multiplexor.

27. A method according to Claim 26, further comprising:

redirecting the set of pointers for processing the transitory data packets for one of the read queue and the write queue for the associated protocol layer.

28. A method according to Claim 26, further comprising:

redirecting the set of pointers for processing the transitory data packets for both the read queue and the write queue for the associated protocol layer.

29. A method according to Claim 25, wherein the network protocol stack is a TCP/IP-compliant protocol stack, further comprising:

defining a set of TCP/IP-compliant protocol layers, selected from the group comprising at least:

a data link protocol layer;

an Internet (IP) protocol layer;

an Transmission Control Protocol (TCP) layer; and

a User Datagram Protocol (UDP) layer.

SYSTEM AND METHOD FOR NETWORK INTRUSION DETECTION ANALYSIS WITH DIRECT PROTOCOL STACK MONITORING

Abstract

A system and method for detecting network intrusions using a protocol stack multiplexor is described. A network protocol stack includes a plurality of hierarchically structured protocol layers. Each such protocol layer includes a read queue and a write queue for staging transitory data packets and a set of procedures for processing the transitory data packets in accordance with the associated protocol. A protocol stack multiplexor is interfaced directly to at least one such protocol layer through a set of redirected pointers to the processing procedures of the interfaced protocol layer. A data packet collector references at least one of the read queue and the write queue for the associated protocol layer. A data packet exchanger communicates a memory reference to each transitory data packet from the referenced at least one of the read queue and the write queue for the associated protocol layer. An analysis module receives the communicated memory reference and performs intrusion detection based thereon.

Figure 1.

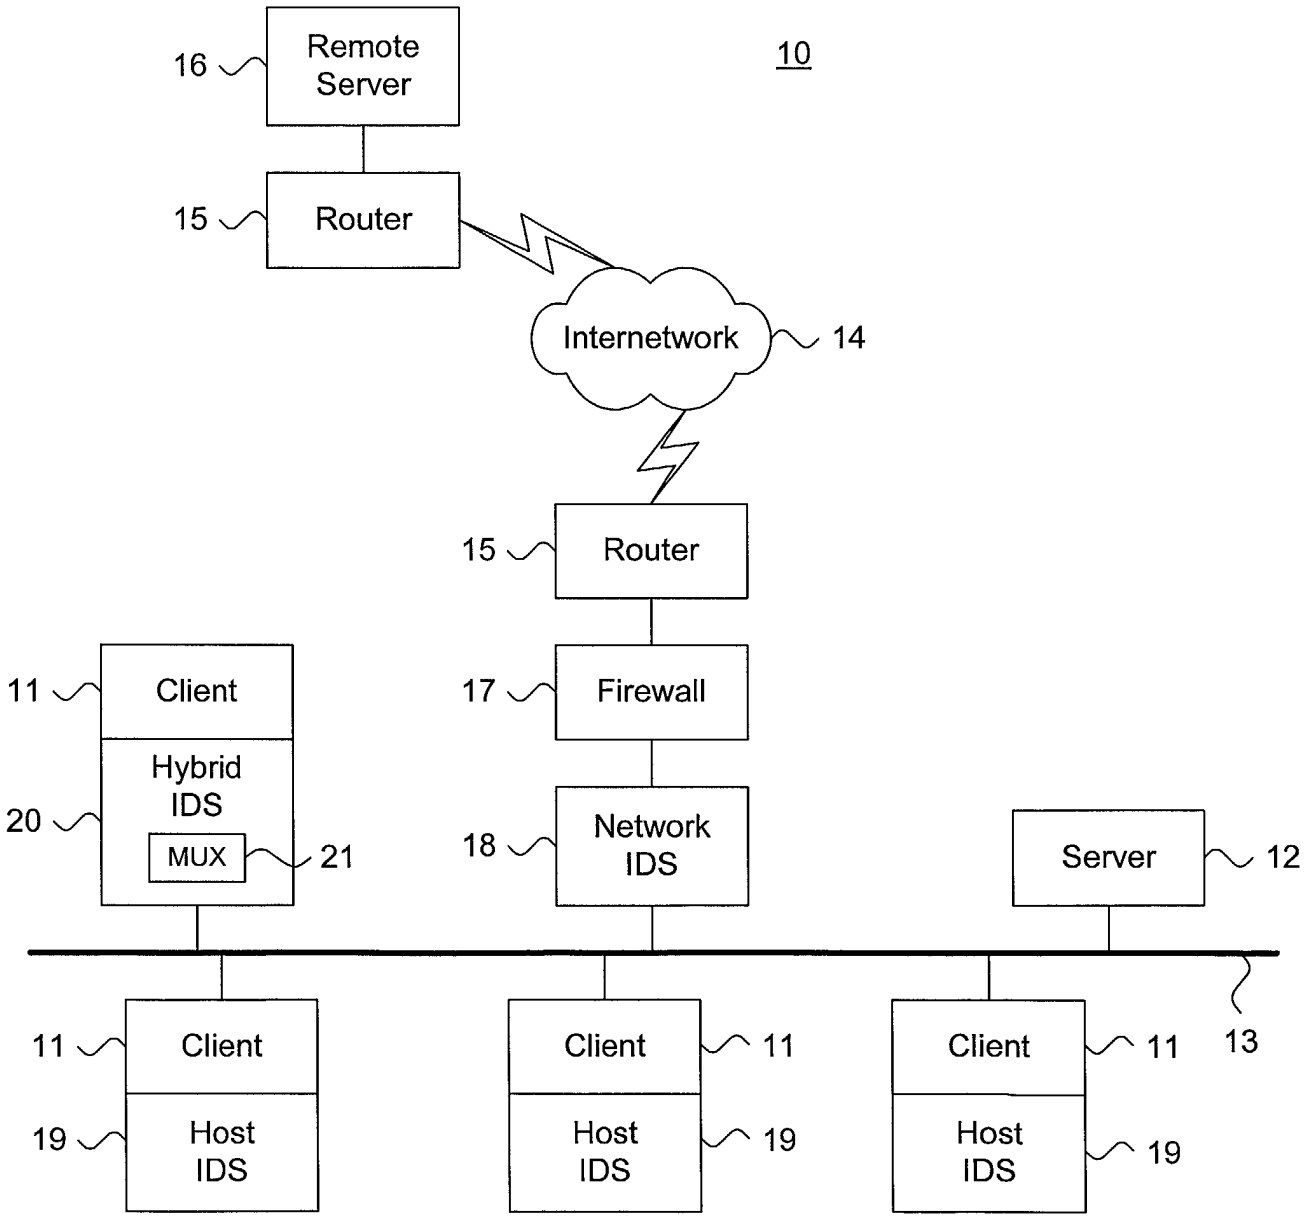

Figure 2 (Prior Art).

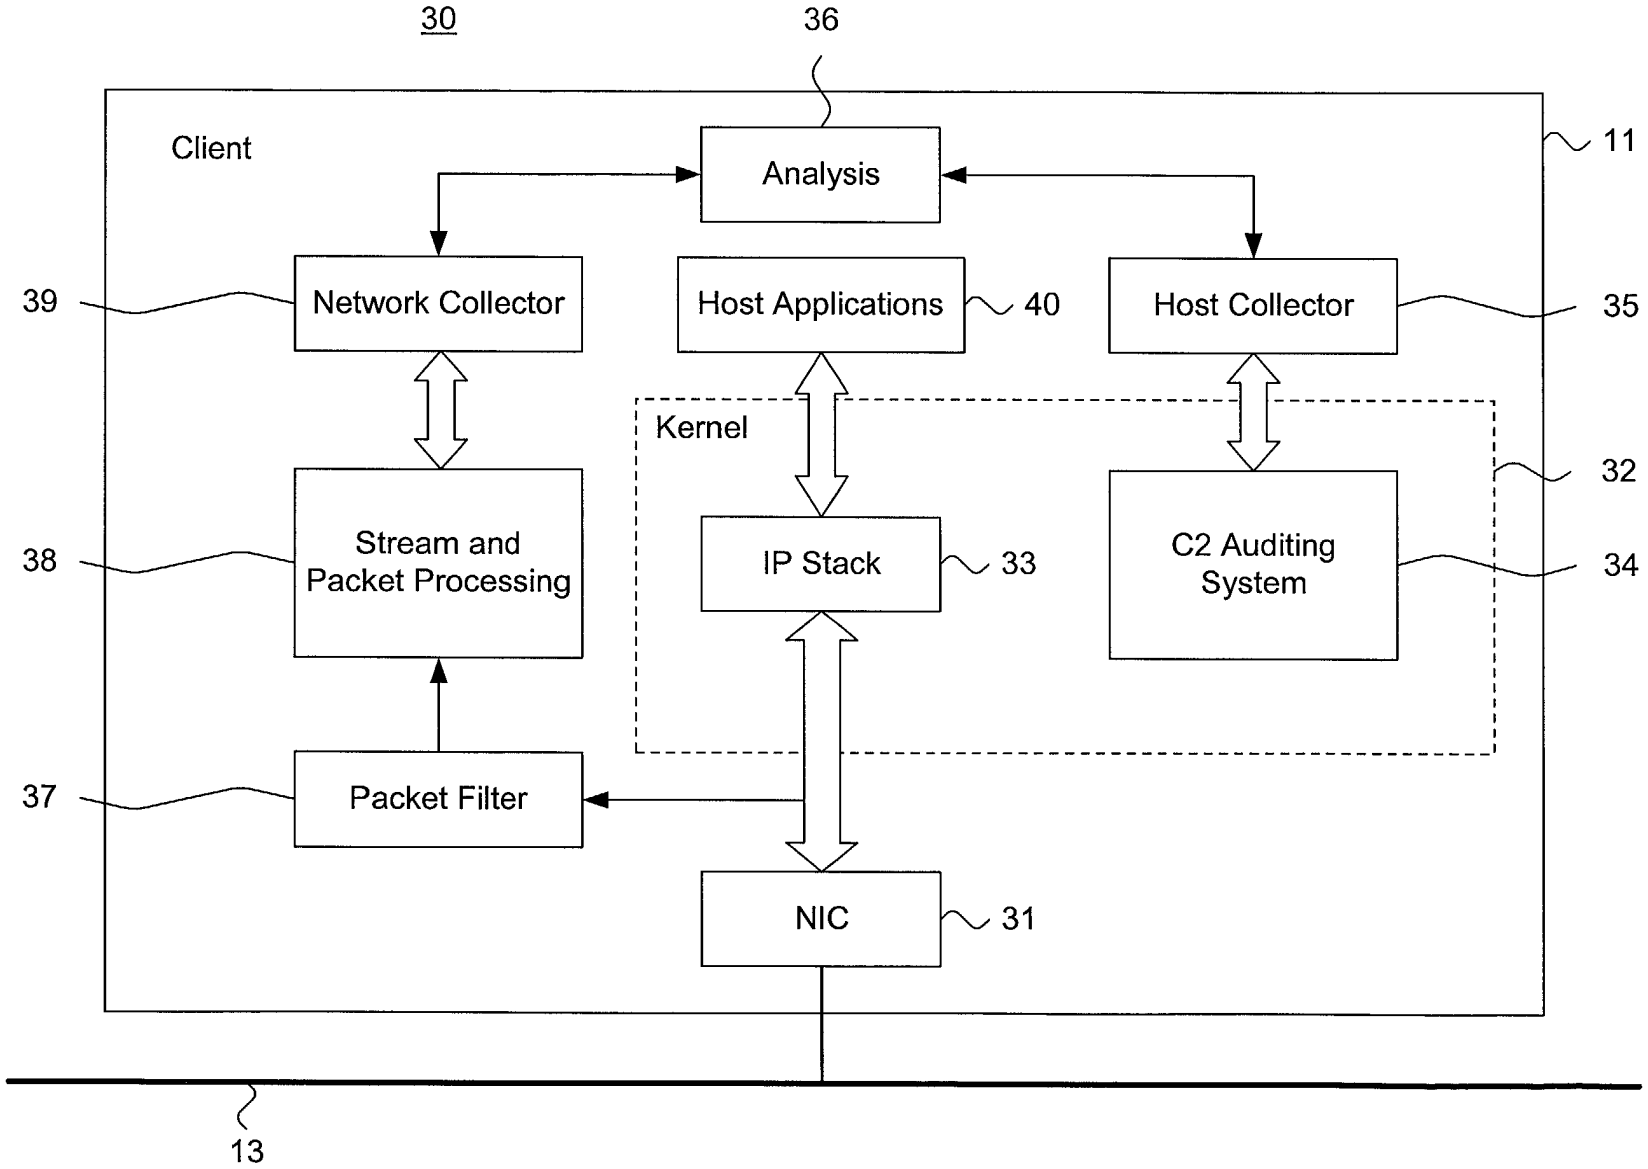

Figure 3.

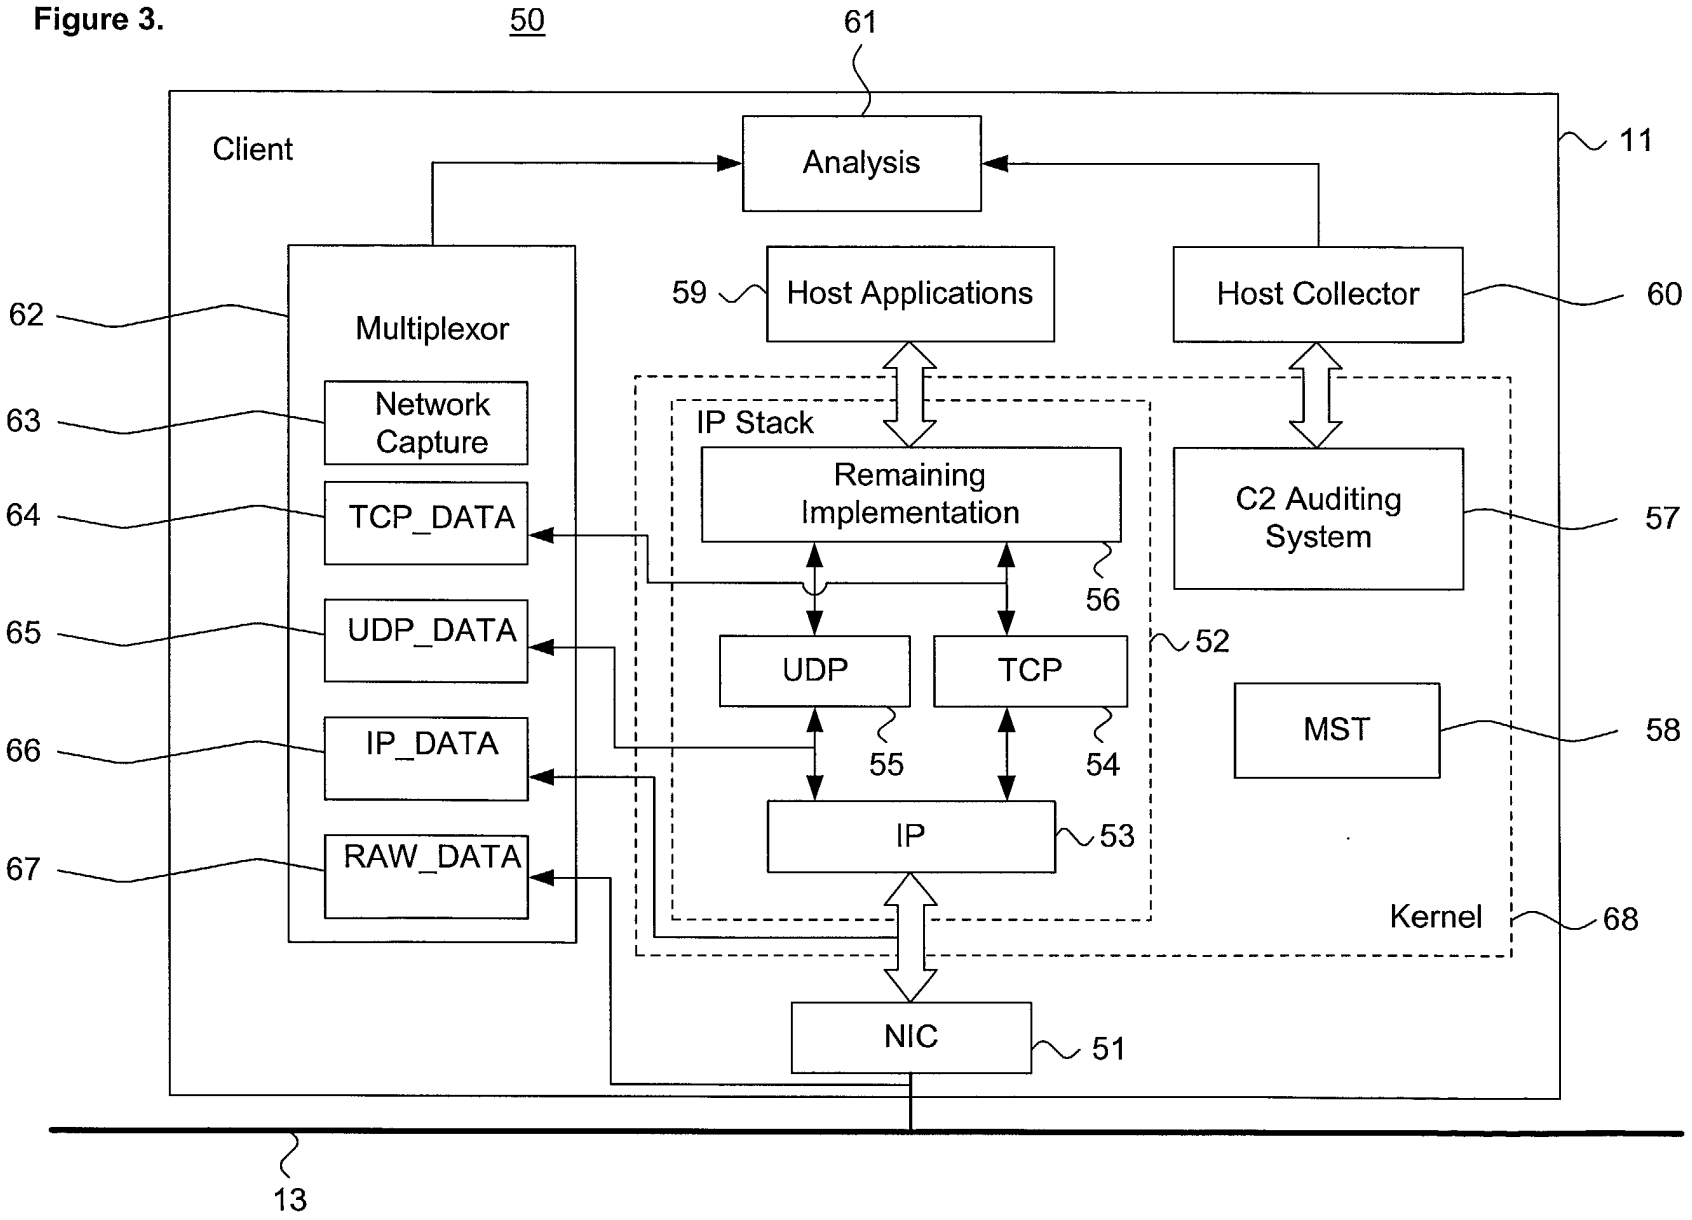

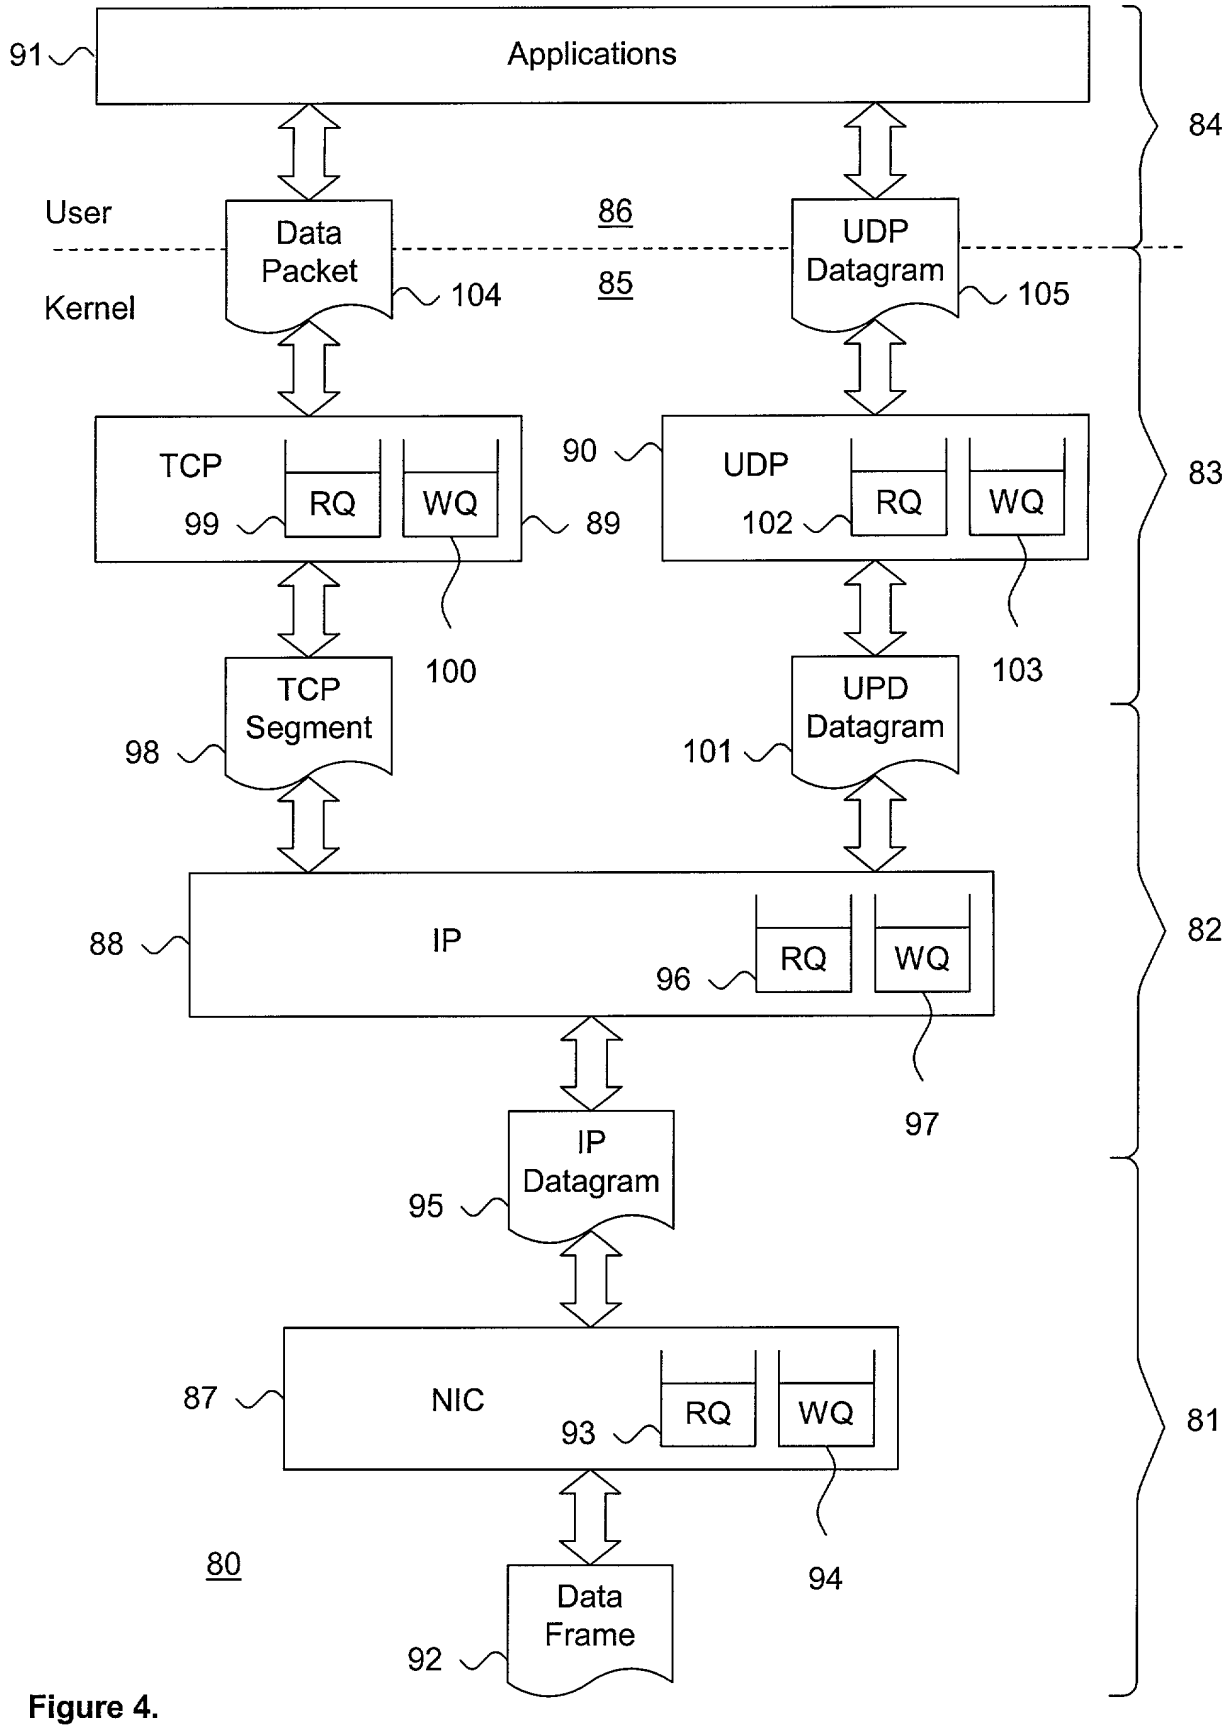
91
Applications
84
User
86
Data Packet
UDP Datagram
Kernel
85
104
105
TCP
RQ
WQ
99
89
90
UDP
102
83
100
103
TCP Segment
UPD Datagram
98
101
88
IP
96
82
97
IP Datagram
95
87
NIC
93
81
80
94
Data Frame
92

Figure 4.

Figure 5.

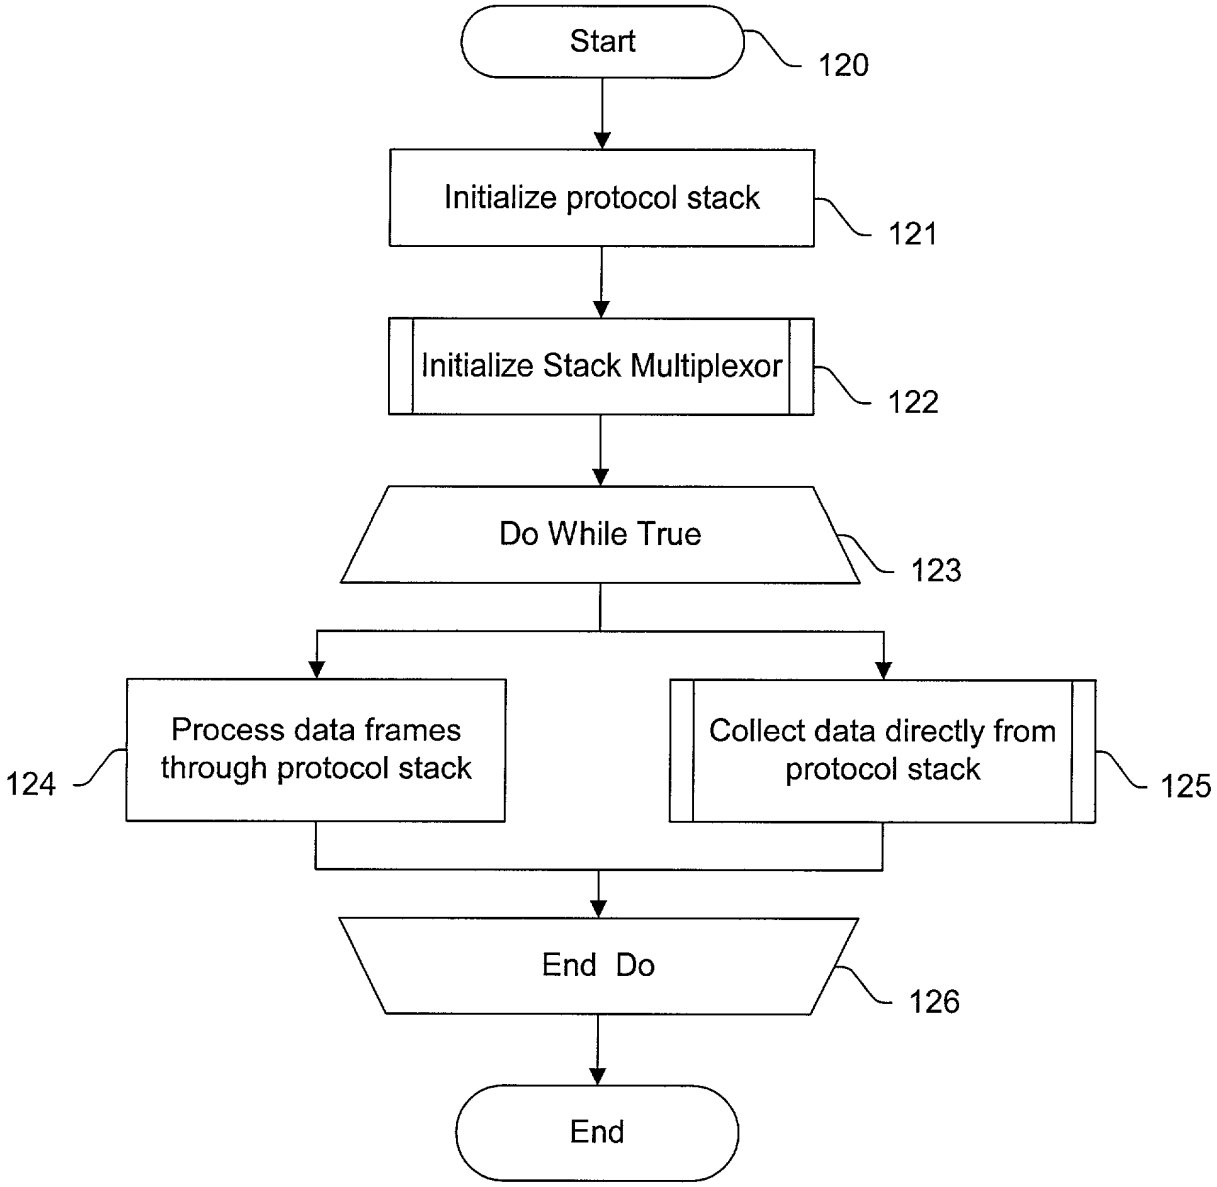
Start
120
Initialize protocol stack
121
Initialize Stack Multiplexor
122
Do While True
123
Process data frames through protocol stack
124
Collect data directly from protocol stack
125
End Do
126
End

Figure 6.

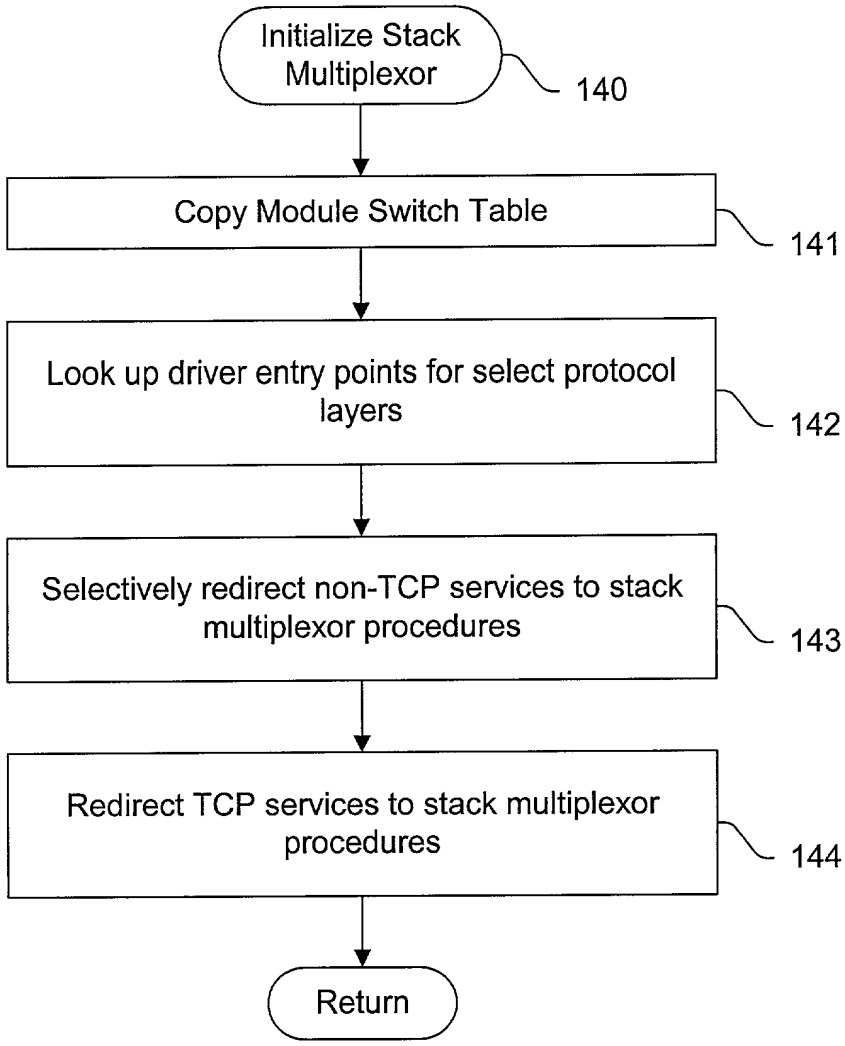
Initialize Stack Multiplexor
140
Copy Module Switch Table
141
Look up driver entry points for select protocol layers
142
Selectively redirect non-TCP services to stack multiplexor procedures
143
Redirect TCP services to stack multiplexor procedures
144
Return

Figure 7.

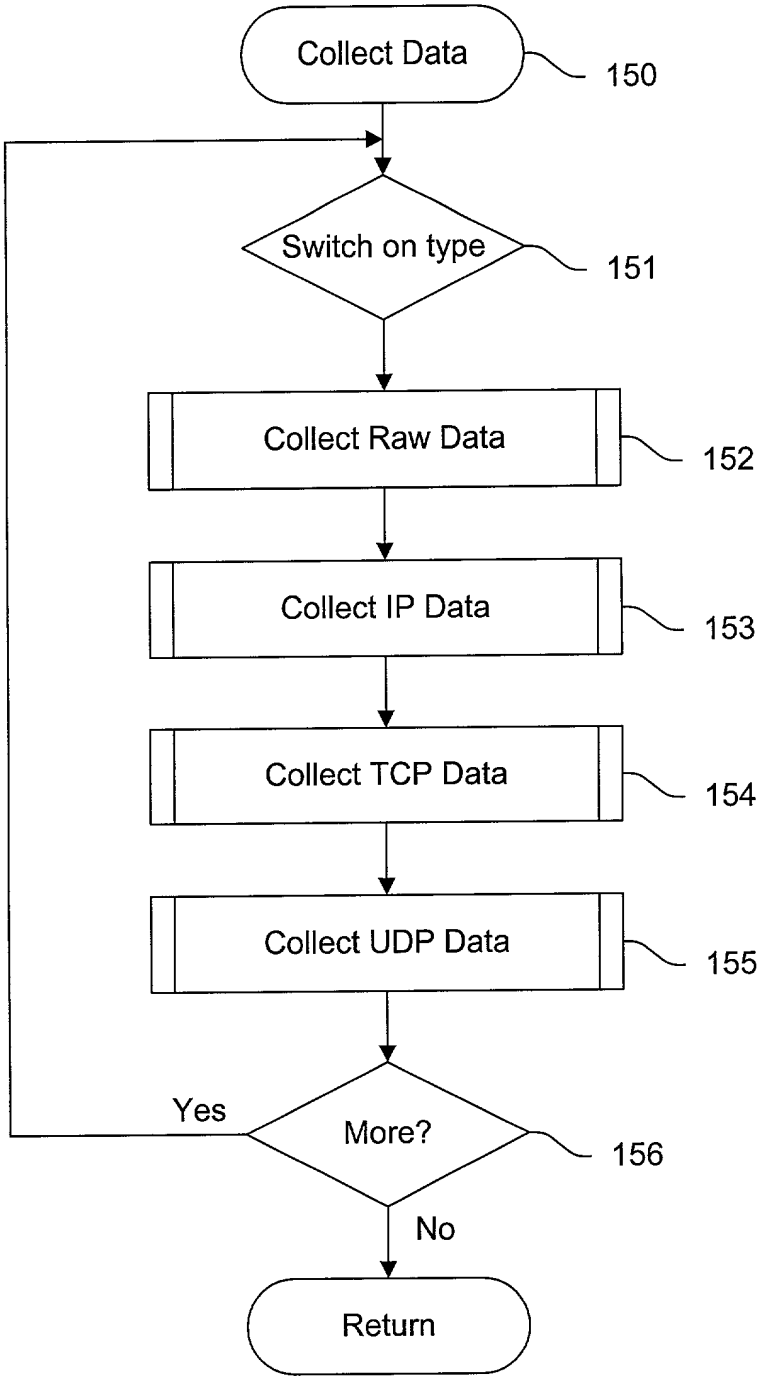
Collect Data
150
Switch on type
151
Collect Raw Data
152
Collect IP Data
153
Collect TCP Data
154
Collect UDP Data
155
Yes
More?
156
No
Return

Figure 8.

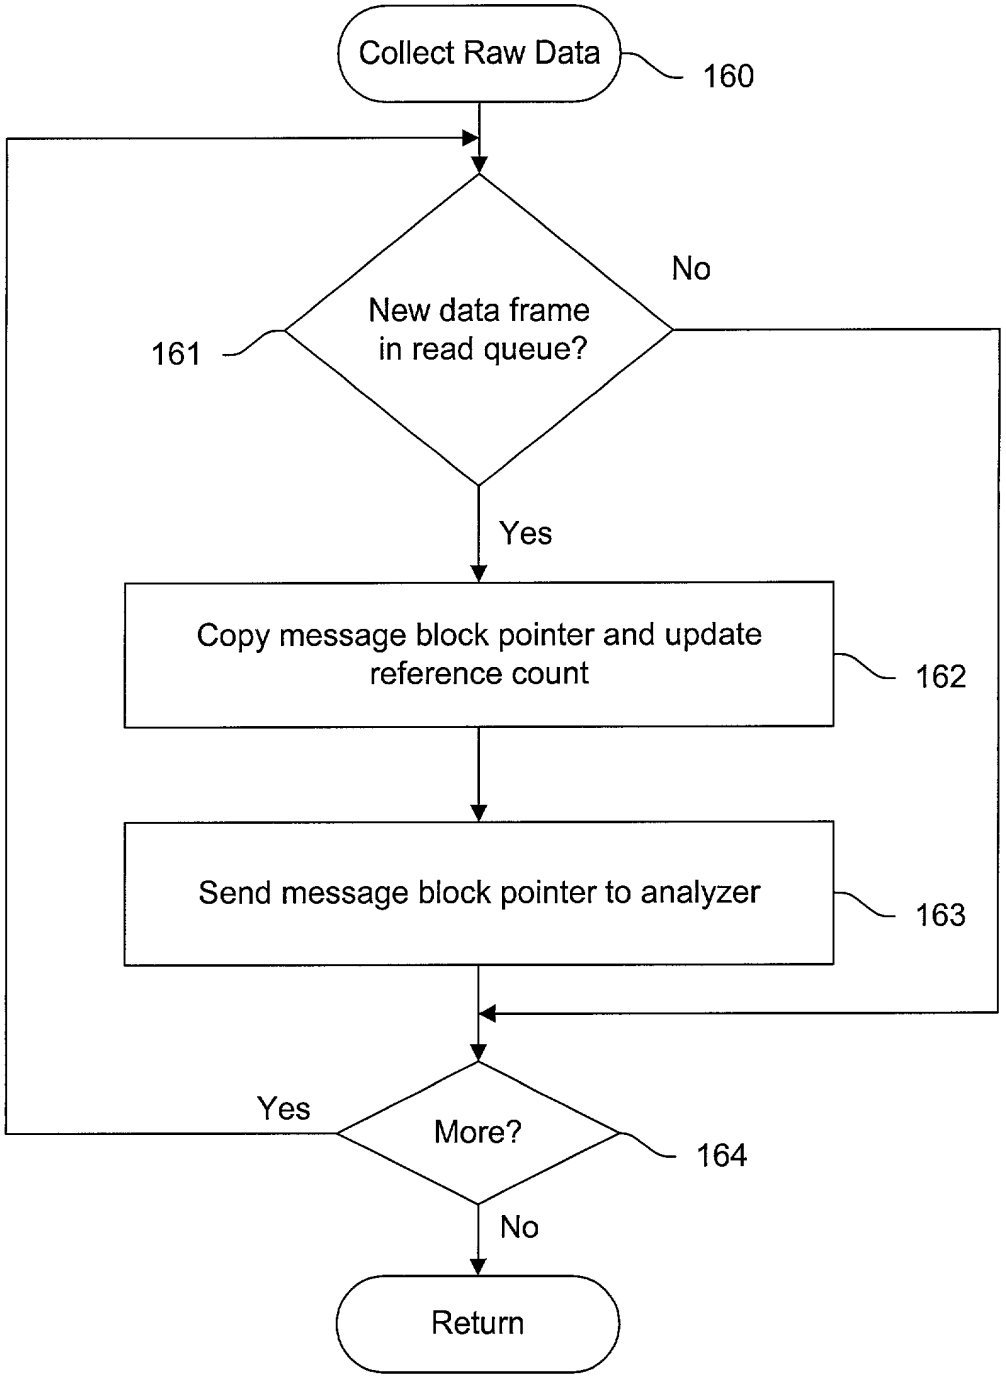
Collect Raw Data
160
New data frame in read queue?
161
No
Yes
Copy message block pointer and update reference count
162
Send message block pointer to analyzer
163
More?
164
Yes
No
Return

Figure 9.

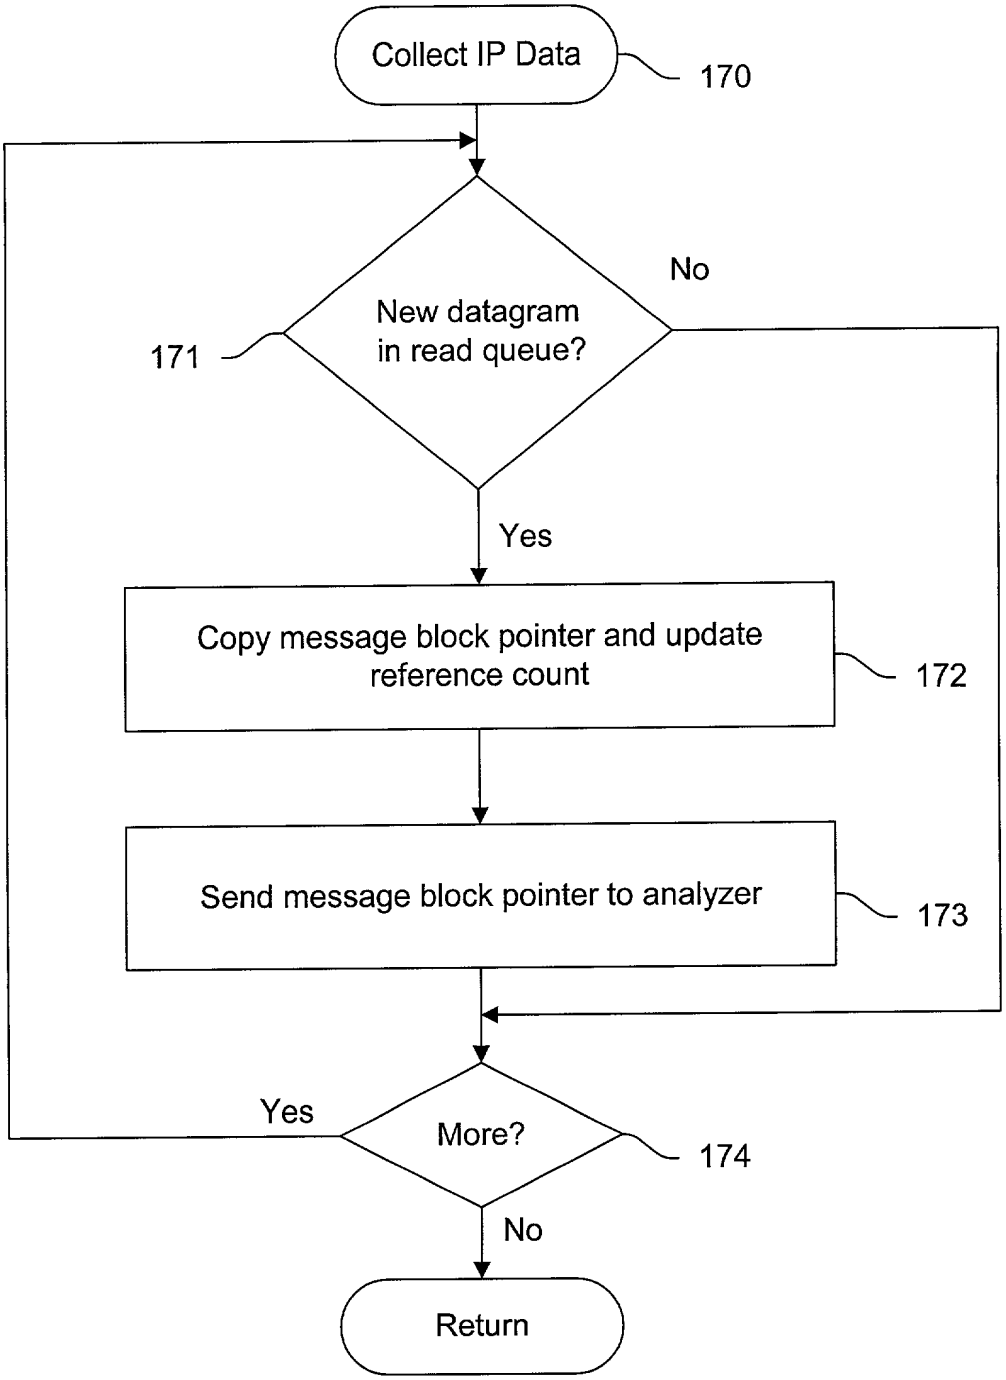
Collect IP Data
170
New datagram in read queue?
171
No
Yes
Copy message block pointer and update reference count
172
Send message block pointer to analyzer
173
More?
174
Yes
No
Return

Figure 10.

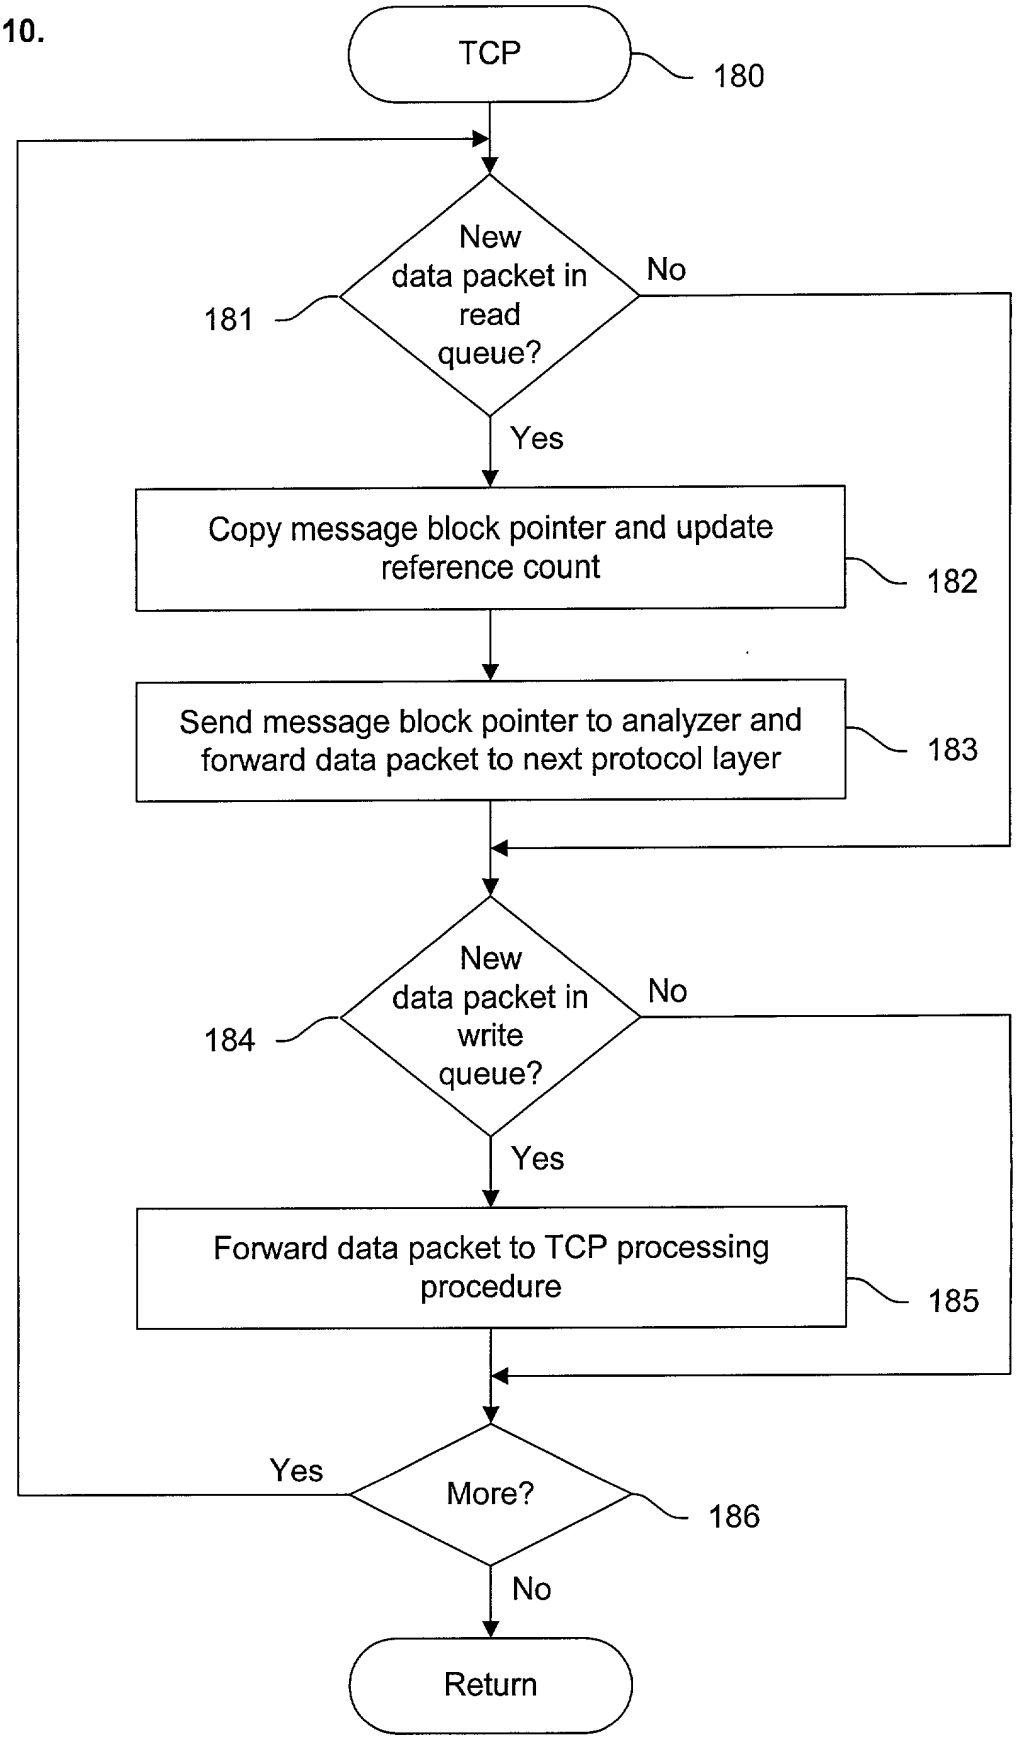
TCP
180
New data packet in read queue?
181
No
Yes
Copy message block pointer and update reference count
182
Send message block pointer to analyzer and forward data packet to next protocol layer
183
New data packet in write queue?
184
No
Yes
Forward data packet to TCP processing procedure
185
More?
186
Yes
No
Return

Figure 11.

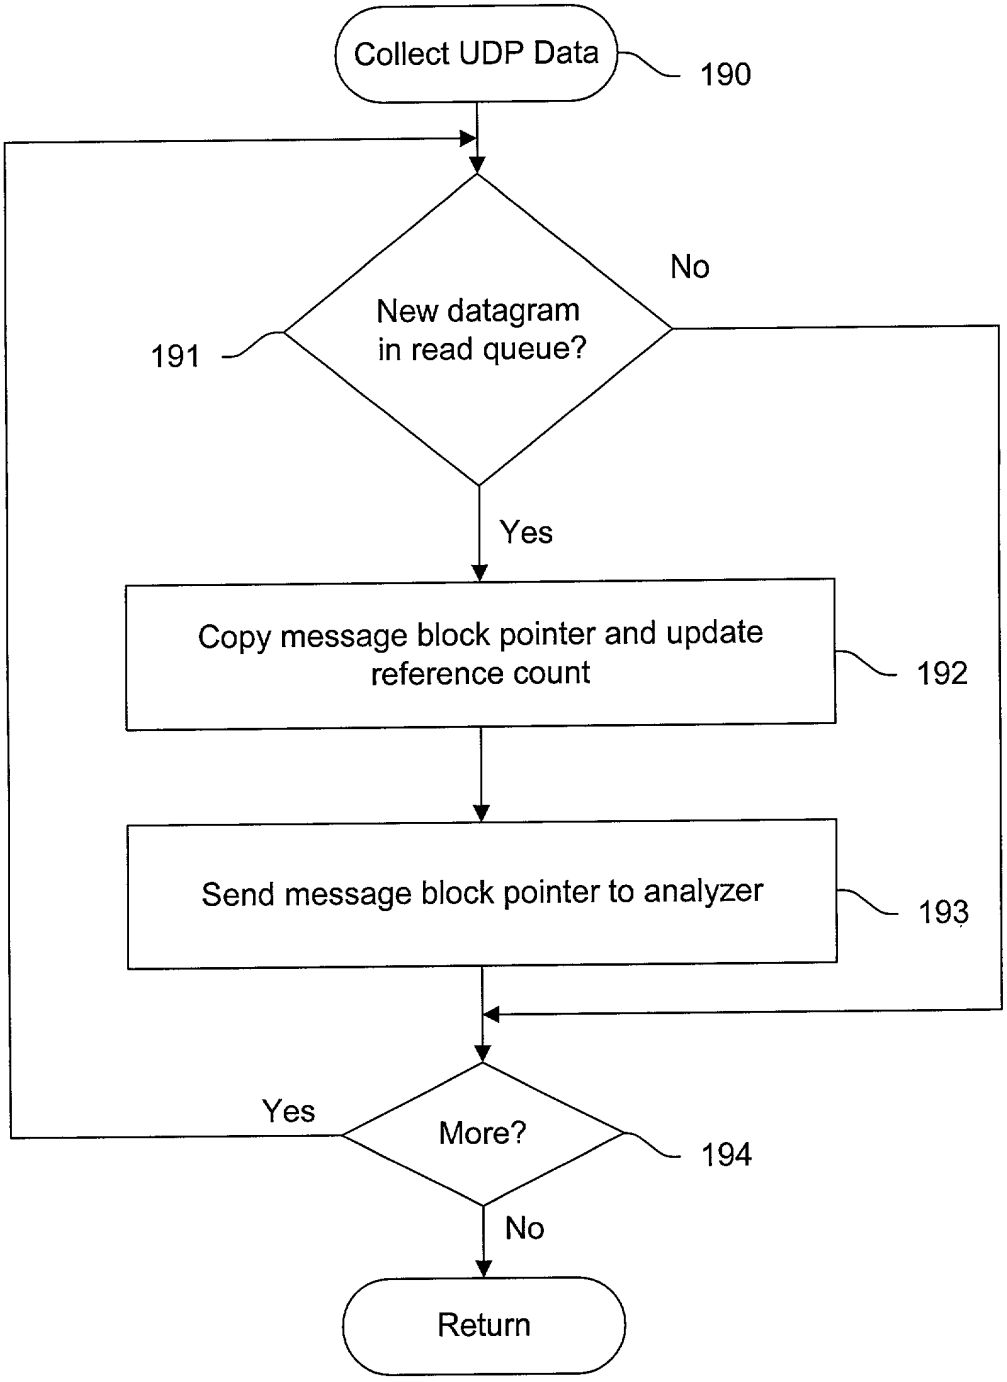
Collect UDP Data
190
New datagram in read queue?
191
No
Yes
Copy message block pointer and update reference count
192
Send message block pointer to analyzer
193
More?
194
Yes
No
Return

PATENT APPLICATION

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION	ATTORNEY DOCKET NO. 002.0141.01

As a below named inventor, I hereby declare that:

My residence/post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

System And Method For Intrusion Detection Data Collection Using A Network Protocol Stack Multiplexor

the specification of which is attached hereto unless the following box is checked:

( ) was filed on _____ as US Application Serial No. or PCT International Application Number _____ and was amended on _____ (if applicable).

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment(s) referred to above. I acknowledge the duty to disclose all information which is material to patentability as defined in 37 CFR 1.56.

Foreign Application(s) and/or Claim of Foreign Priority

I hereby claim foreign priority benefits under Title 35, United States Code Section 119 of any foreign application(s) for patent or inventor(s) certificate listed below and have also identified below any foreign application for patent or inventor(s) certificate having a filing date before that of the application on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE FILED	PRIORITY CLAIMED UNDER 35 U.S.C. 119
			YES:___ NO:___
			YES:___ NO:___

Provisional Application

I hereby claim the benefit under Title 35, United States Code Section 119(e) of any United States provisional application(s) listed below:

APPLICATION SERIAL NUMBER	FILING DATE
60/182,842	2/16/2000

U.S. Priority Claim

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

APPLICATION SERIAL NUMBER	FILING DATE	STATUS(patented/pending/abandoned)

POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) listed below to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

Patrick J.S. Inouye, Esq, Reg. No. 40297 **Brian J. Daiuto, Esq, Reg. No. 38422**

Send Correspondence to:	**Direct Telephone Calls To:**
Patrick J.S. Inouye, Esq **Patrick J.S. Inouye, P.S.** **816 Second Avenue** **P.O. Box 21808** **Seattle, WA 98111-3808**	**Patrick J.S. Inouye, Esq** **(206) 381-3900**

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Inventor: Daniel T. Holland III **Citizenship:** USA

Residence: 358 Colville Drive, San Jose, California 95123-3506

Post Office Address: Same

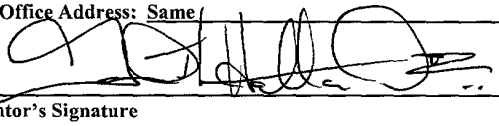

Inventor's Signature **Date** 9 Aug 2000

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION (continued)

ATTORNEY DOCKET NO. 002.0141.01

Full Name of Inventor: Roark B. Hilomen

Citizenship: USA

Residence: 1270 Peiking Drive, San Jose, California 95131

Post Office Address: Same

Inventor's Signature

Aug 9, 2000

Date

Full Name of Inventor: Steven P. Lang

Citizenship: USA

Residence: 777 W. Middlefield Road #135, Mountain View, California 94043

Post Office Address: Same

Inventor's Signature

Date

POWER OF ATTORNEY BY ASSIGNEE TO EXCLUSION OF INVENTOR UNDER 37 C.F.R. § 3.71 WITH REVOCATION OF PRIOR POWERS

Applicant(s): Holland et al.
Title: System And Method For Intrusion Detection Data Collection Using A Network Protocol Stack Multiplexor
Filing Date: TBD
Serial No.: TBD

Group Art Unit: *Unassigned*
Examiner: *Unassigned*
Attorney Docket No: 002.0141.01

The undersigned ASSIGNEE of the entire interest in the above-identified application for letters patent hereby appoints Brian J. Daiuto, Reg. No. 38,422 of NETWORK ASSOCIATES, INC., and Patrick J.S. Inouye, Reg. No. 40,297 of PATRICK J.S. INOUYE, P.S., to prosecute this application and transact all business in the United States Patent and Trademark Office in connection therewith and hereby revokes all prior powers of attorney; said appointment to be to the exclusion of the inventors and the inventors' attorneys in accordance with the provisions of 37 C.F.R. § 3.71.

The following evidentiary documents establish a chain of title from the original owner(s) to the Assignee:

_X_ a copy of an Assignment attached hereto, which Assignment has been (or is herewith) forwarded to the Patent and Trademark Office for recording; or

___ the Assignment recorded on _____ at reel _____, frame ____.

Pursuant to 37 C.F.R.§ 3.73(b) the undersigned Assignee hereby states that evidentiary documents have been reviewed and hereby certifies that, to the best of ASSIGNEE's knowledge and belief, title is in the identified ASSIGNEE.

Please direct all telephone calls and correspondence to: Patrick J.S. Inouye, Patrick J.S. Inouye, P.S., P.O. Box 21808, Seattle, WA 98111-3808; telephone: (206) 381-3900; facsimile: (206) 381-3999.

ASSIGNEE: Networks Associates, Inc. d/b/a Network Associates, Inc.

Signature: _____ 7/31/02
(Signature) (Date)

Name: Brian J. Daiuto, Esq.

Title: Senior Manager of Intellectual Property